Modeling Adaptive Learning: R&D Strategies in the Model of Nelson & Winter (1982)

Murat Yıldızoğlu
IFREDE-E3i
Université Montesquieu Bordeaux IV
Avenue Léon Duguit
Fax. 335 56 84 86 47
e-mail: yildi@montesquieu.u-bordeaux.fr
http://yildizoglu.montesquieu.u-bordeaux.fr/

May 2001 – Preliminary version v.0.5

Abstract

This article aims to test the relevance of learning through Genetic Algorithms **(GA)** and Learning Classifier Systems **(LCS)**, in opposition with fixed R&D rules, in a simplified version of the evolutionary industry model of Nelson and Winter. These three R&D strategies are compared from the points of view of industry performance (welfare): the results of simulations clearly show that learning is a source of technological and social efficiency.
Keywords: Innovation, Industry dynamics, Bounded rationality, Learning, Genetic algorithms, Learning classifier systems
JEL: O3, L1, D83, D84

1 Introduction

Research and development (R&D) decisions are characterized by a strong uncertainty concerning the return on investment. This uncertainty is stronger for R&D investment than for other types of investment. Indeed innovations often result from what Simon [1958] calls "nonprogrammed decisions", that is situations where the alternatives of choices must be discovered by firms and the connections between choices and consequences are imperfectly known. It is the reason why R&D decisions are generally associated with the uncertainty in the sense of Knight [1921]. This uncertainty strongly limits the ability of firms to form expectations about the return on their R&D investment. In this context, firms must be able to improve, through experience, their perception of the relationships between R&D investment and competitiveness and to adapt accordingly their R&D decisions (see Oltra & Yildizoglu [1999]). Modeling R&D decisions must consequently rely on decision rules more sophisticated than the fixed rules that have been traditionally adopted in models of technology dynamics under bounded rationality assumption (see, for example, the models in Nelson & Winter [1982]).

In fact, fixed rules are not a necessary characteristic of decisions under bounded rationality: the main characteristic is procedural rationality (Simon [1976]). Hence, bounded rationality does not preclude the tendency of agents to adjust their decisions to the evolution of their (technological and competitive) environment. Even if they do not search for the globally best solutions, agents learn from their experience and this learning allows them to *fine-tune* their decisions. Consequently, one must search for a better way of modeling decisions in order to take into account this individual learning.

Of course, one could choose to implement a simple process of trial and error but such a process would contain a strong ad hoc element in the way it models the sequence of trials. Firms do definitely not proceed by purely random trials. When learning is individual[1], new strategies are necessarily based on the past experience: firms combine known decision rules in order to reach better ones. Genetic Algorithms (**GA**), implement such a learning process through evolutionary mechanisms: from a population of actual decision rules, the selection keeps the best ones, the crossover combines these in order to obtain better rules and the mutation introduces some small amount of random experimenting. Since they take into account learning, they constitute a good rival of fixed rules that seem to actually constitute the only unifying approach.

Yildizoglu [2001] tests the relevance of GA, in opposition with fixed R&D rules, in a simplified version of the evolutionary industry model of Nelson and Winter. The original model is simplified in order to focus on R&D process as the main determinant of industry dynamics. Firms arbitrate at each period between R&D and physical capital when allocating their gross profits to different investment projects. The industry is composed of two types of firms: the **NWFirms** that use a fixed decision rule and **GenFirms** that adjust their R&D investment through a GA. Competition selects, in the long term, the firms that outperforms their competitors: firms can only finance their investments by the profits and they must leave the industry when their physical capital vanishes. This article shows that at the industry level and from the point of view of competitive power, learning represented by the GA is a source of efficiency.

But this model relies on a very restrictive assumption about the learning process of the firms because it uses a very simple implementation GAs: firms effectively use each rule in order to evaluate its fitness. This view is very unsatisfactory given that the experience of the firms will generally be used by them to form expectations on the workings of their environment. To include these expectations firms must be able to generalize from their experience. Learning Classifier Systems (**LCS**) implement such generalizations. As Pier Luca Lanzi very nicely summarizes it (see Lanzi, Stoltzmann & Wilson [2000], p. 20):

> In learning classifier systems an agent learns to perform a certain task by interacting with a partially unknown environment from which the agent receives feedback in the form of numerical reward. The incoming reward is exploited to guide the evolution of the agent's behavior which, in learning classifier systems, is represented by a set of rules, the classifiers. In particular, temporal difference learning is used to estimate the goodness of classifiers in terms of future reward; genetic algorithms are used to favour the reproduction and recombination of better classifiers.

This is why this new article builds upon Yildizoglu [2001] and, in order to evaluate the relevance of LCSs, it includes a third type of firms in the model: **XCSFirms**. XCSFirms use a particular implementation of LCS proposed by Wilson [1995].

Our results show that industries formed by XCSFirms are even more efficient than the ones composed by the GenFirms. The use of XCSs for modeling learning emphasizes the role of learning in the efficiency of industries. These results are obtained through the simulation methodology already developed in Jonard & Yildizoglu [1998]. This procedure uses the comparison, with non-parametric statistical methods, of the results of batches of simulations instead of the comparison of individual simulations. In order to better understand the learning process we scrutinize the last simulation in each batch.

A jar file containing the Java 2 version of the class files of the model can be found on our web site. This site also contains full documentation in Sun's API format.

The remainder of this article is organized as follows. In section two we present the characteristics of the model. The connection between the GAs, XCSs and the learning process also is discussed it this section. Section three is dedicated to the presentation of our methodology and results. Section four concludes.

2 The model

I only emphasize new elements in the model. The intersection with the well known Nelson & Winter [1982] (part V, ch.12) model will first be outlined. A second section will present new dimensions included in this

[1] Silverberg & Verspagen [1996] as well as Kwasnicki & Kwasnicka [1992] formalize learning as local individual random experiments combined with industry wide imitation. This learning process has both an individual and a population level components. Our objective is to model learning as a fully individual process.

model: capital and R&D investments.

2.1 Characteristics common with Nelson and Winter(1982)

At the beginning of each period, the firm j is characterized by the productivity A_j of it's technology and it's capital stock, K_j. Capital is the only production factor, and the production technology is characterized by fixed input coefficient and constant scale economies. Unit using cost of capital, c, is constant over different production techniques (the unit cost of production is c/A_j). The capital stock depreciates at a rate δ at each period.

Production technics are disembodied. There is no switching cost and the capital can be converted without cost from one technology to another (for a more realistic model with vintage capital, see Silverberg, Dosi & Orsenigo [1988]). This corresponds to a vision of technology based on process innovation. In fact, the innovating firm does not replace its capital stock, but uses it more efficiently. An innovation therefore corresponds to better knowledge of the production process.

2.1.1 Production and profits

Each firm in the industry ($j \in I = \{1...N\}$) produces the same homogenous good with the following production function:

$$Q_j = A_j \cdot K_j. \tag{1}$$

The gross profit rate on capital of the firm is given by:

$$\pi_j = pA - c \tag{2}$$

where p is the market price and is determined by a short term equilibrium on the product market:

$$\begin{cases} Q = \sum\limits_j Q_j \\ p = p(Q) = \dfrac{\mathbf{D}}{Q^{1/\eta}} \end{cases} \tag{3}$$

where Q is the total supply, $p(Q)$ is a constant elasticity inverse market demand function, and η is the Marshalian demand elasticity. Gross profits of the firm are given by

$$\Pi_j = \pi_j \cdot K_j \tag{4}$$

The state of each firm will change from one period to another in consequence of the R&D decisions, which modify its technology and hence its productivity, and the investment behavior, which modifies its capital stock.

2.1.2 R&D and technical progress

The productivities are modified in each period consequently to the technical progress. In each period firms invest RD_{jt} on R&D. This investment allows them to imitate their successful competitors and to innovate. Both imitation and innovation are two-stage stochastic processes.

Innovation

Innovation is a two-stage stochastic process. A first draw determines if the R&D investment of the firm has been successful and resulted in an innovation:

$$P[d_{int} = 1] = a_n \cdot RD_{jt}, \tag{5}$$

where a_n is a calibration parameter that projects RD on $[0,1]$. A second draw gives the effective result of the innovation

$$\tilde{A}_{jt} \rightsquigarrow N\left(A_{jt}, \sigma_2^2\right).$$

Hence innovation is a cumulative process and firms with higher productivities have better chance to attain even higher productivities.

Imitation

For the imitation, we have one stochastic draw which determines if the firm's R&D investment has been successful. If it is the case, the firm obtains the best practice in the industry (A_t^*):

$$P[d_{imt} = 1] = a_m \cdot RD_{jt}$$
$$\hat{A}_{jt} = A_{jt} + d_{imt} \cdot (A_t^* - A_{jt}).$$

New productivity of the firm

Finally, the effective productivity of the firm for the next period is given by the best of these three outcomes:

$$A_{j,t+1} = \max\left\{A_{jt}, \tilde{A}_{jt}, \hat{A}_{jt}\right\} \tag{6}$$

2.2 Capital investment and R&D decisions

Main differences between this model and Nelson & Winter (1982) consist in the investment behavior: investment in physical capital and investment in R&D. A possibility of exit from the industry is also included in the model. In each period firms invest a fraction of their gross profit on R&D. The rest of this profit is used for the expansion of physical capital.

2.2.1 R&D decisions and learning

Firms invest a fraction rd_{jt} of their gross profits on R&D. A minimal investment is necessary to keep *alive* the R&D potential (research equipment and team). We therefore have $rd_{jt} \geq rd_{\min}$.

There are three types of firms: NWFirms, GenFirms and XCSFirms. They are distinguished by their R&D investment behavior.

NWFirms invest in each period a fixed proportion rd_{NW} of their profit in R&D (in addition to the minimal amount of R&D):

$$RD_{NW\,jt} = \max\{rd_{\min}, rd_{NW}\} \cdot \Pi_{jt} \tag{7}$$

This rule corresponds to the representation of bounded rationality by "fixed rules". Learning of firms about their environment does not influence their R&D behavior. This is the common approach retained in many evolutionary industry models. Learning is taken into account in the behavior of GenFirms.

Each **GenFirm** uses an individual genetic algorithm (GA) in order to adjust the R&D strategy (the fraction $rd_{jt} \geq rd_{\min}$) to the conditions of the industry. Each possible strategy of the firm is coded as a chromosome C_i of length **G**. During its life, the firm carries a population of **C** strategies (number of chromosomes). This population of parallel rules evolves as a consequence of the experience of the firm in the industry.

The experience of the firm can only influence these rules if it provides an evaluation mechanism for different rules. In an industrial context, the only way of evaluating a rule is using it: the value of a rule depends on the dynamics of the industry and hence, on the behavior of other firms. Moreover, R&D investment does not pay back immediately and each R&D strategy must be used for many periods before proving its efficiency. Consequently, in order to evaluate each rule, the firm uses it for a number of periods ($n =$*learning period*) and **the average gross profit rate of this time interval gives the fitness of this strategy**. When all strategies of the population are evaluated, a new population is generated through selection–crossover–mutation. We use an elitist GA that conserves the best strategy of the preceding period in the population.

We adopt an indirect coding of R&D strategies: the fraction of profits dedicated to R&D (strategy) is coded as a chromosome C_i of length **G**. The decimal value of the chromosome corresponds to the position of this strategy in the search space $[0\%, 100\%]$. This space contains $\Delta = \sum_{i=0}^{i=G-1} 1 \cdot 2^i$ equally spaced strictly positive strategies, and zero. The R&D rate corresponding to a chromosome C_i is finally computed using the following rule

$$rd_j = (C_i)_{10} \cdot \frac{1}{\Delta}. \tag{8}$$

Example: If $G = 4$, there are $(1111)_{10} = 1 \cdot 2^3 + 1 \cdot 2^2 + 1 \cdot 2^1 + 1 \cdot 2^0 = 15 = \Delta$ strictly positive strategies equally spaced between 0% and 100%. If a strategy of the firm is $C_i = 0011$, this chromosome corresponds to the following R&D investment rate:

$$C_i = (0011)_{10} = 3 \Rightarrow rd_j = 3 \cdot \frac{1}{15} = 20\%.$$

Finally, the R&D investment of the firm is given by

$$RD_{jt} = \max\left\{rd_{\min}, rd_{jt}\right\} \cdot \Pi_{jt} \tag{9}$$

Even if the GA does not represent the exact learning mechanism of firms, it is a convenient way of representing the presence of this learning at the individual level. A more detailed discussion of the connection between the learning process of the firms and the GA can be found in Yildizoglu [2001]. But it should be clear that such a use of GA limits the learning of the firms to the exploration of their strategy space (R&D rates over profit). These strategies are *adapted* to the evolution of the environment but the firm does not determine satisfactory strategies for different states of the environment (it does not really possess a mental model of the environment – see Oltra & Yildizoglu [1999]). This is why it has to effectively implement each strategy in order to evaluate it.

Learning following this stronger sens can only be present if firms can *generalize* from their past experience in order to anticipate the value of each considered strategy before using it. This possibility is included in the model through the XCSFirms.

XCSFirms LCSs, proposed first by John Holland, allow for such a generalisation by combining decision rules with a *condition* part that characterizes the states of the environment for which that decision can be relevant. As a consequence, the firm does not have to consider all strategies in all situations (states of the environment) but only the relevant strategies.

XCS, proposed by Wilson [1995], implements a very satisfactory version of Holland's ideas mainly by defining the fitness of a rule in terms of its accuracy (instead of its strength), by eliminating the internal message system and by using a niche based GA for the evolution of decision rules.

An XCS is composed by a population of classifiers composed by a condition part and a rule part. The condition part is composed from an alphabet with three letters $\{0, 1, \#\}$. The last letter should be read "*don't mind*" and it is compatible with either 0 or 1. Following a signal received from the environment (the *state*), a **match set** (M) is formed from the population (P) using all classifiers with the condition part that matches this signal (see Figure 1). The relative fitness of the rules determines the action that should be chosen from this set on the base of the accuracies of all included classifiers. All classifiers prescribing this action form the **action set** (A). The implementation of this action yields a **reward** from the environment and this reward is used for updating the fitnesses of classifiers included in the action set. The exploration component is a GA that modifies the match set (niche based GA). Consequently, the XCS is composed of three blocks:

- The *performance component* that is in charge the choice of the action that will be implemented (from P, through the sets M and A):
- The *reinforcement component* that is in charge of the update of the fitness of the elements of A;

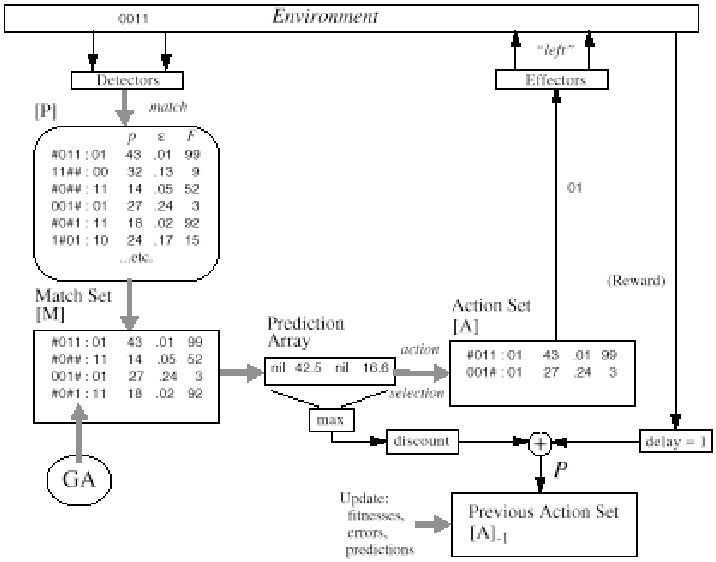

Figure 1: XCS (Source: Wilson(1995))

- The *discovery component*: the GA on M.

A more detailed description can be found in Wilson [1995].

Through the # letter a classifier can prescribe its action for different states of the environment. This is the base of generalization in LCS. But in multistage problems (the investment profile of a firm living in an industry is such a problem), this can be the source of over-generalization: the population can be dominated by general rules that will necessarily be quite bad for at least some situations. The innovations included in XCS have revived the LCS research program because they favour a better understanding of this mechanism. Using niche-based exploration, XCS seriously limits over-generalisation and allows for the emergence of performant specific rules. As such, they provide a better trade off between the generalization, which is necessary for representing learning of firms, and the performance, which limits the excessive cost of the evolutionary representation of learning.
Each **XCSFirm** in our model explores the space of R&D investment strategies $(actions : rd \in [0,1])$ through its XCS (we use a modified version of the Java version proposed by Butz & Wilson [2000] and Butz [2000]).

The detectors of the XCS are coded on eight bit and they can detect if

- the market price has increased or not;
- the concentration of the capital has increased or not;
- the average productivity has increased or not;
- the number of active firms has decreased or not;
- the profit of the firm has increased or not;
- the productivity of the firm has increased or not;
- the capital stock of the firm has increased or not;
- R&D investment volume of the firm has increased or not.

We explore two possibilities for the reward of the actions: the profit rate and the productivity of the capital. Other parameters of the XCS we use in our simulations are given in the appendix.

2.2.2 Investment in physical capital

Capital investment results directly from the arbitrage of firms between the R&D investment and capital expansion. Learning firms adapt the sharing of gross profits between R&D and physical capital:

$$K_{jt+1} = (1-\delta)\cdot K_{jt} + \left(1-\max\left\{rd_{\min}, rd_{jt}\right\}\right)\cdot \Pi_{jt}. \tag{10}$$

2.2.3 Exit

If the profits of the firm get persistently low, it can loose all possibility of investment and innovation. In this case, current profits do not permit any investments. The capital stock of the firm vanishes because of the depreciation. When the capital stock gets very small, the firm looses all possibility of innovation and growth. It consequently exits the industry when

$$\pi_{jt} < 0 \text{ and } K_{jt} \leq \underline{K}.$$

3 Simulation methodology and results

We use the simulation protocol developed in Jonard & Yildizoglu [1998]. This protocol is explained in a first paragraph. Relative performance of GenFirms is measured through different indicators that have been developed for this article. Simulation results are used to assess the influence of learning on the aggregate performance of the industry.

3.1 Protocol

Since we aim to derive results independent from a particular sequence of random numbers, a batch of 20 simulations, of 2000 periods each, is run for each configuration of the model. Observations have been saved every 40 periods. The whole possible history of the industry is hence represented by a sample of 1000 observations. The relevant dimensions (e.g. technical progress, concentration) of resulting samples are compared by way of non-parametric testing (the non-parametric Wilcoxon–Mann–Whitney test, see for instance ch.18 in Watson, Billingsley, Croft & Huntsberger [1993]). For convenience, results are presented as **box plots** where the box gives the central 50% of the sample centered around the median: the box hence gives the first, second and third quartiles (Q_1, Q_2, Q_3) of the distribution. The whiskers give the significant minimum and the significant maximum of the distribution. Each box contains the whole history of the industry for all simulations for each corresponding configuration.

This protocol allows the qualitative comparison of different industry configurations. Different indicators are used for these comparisons.

3.2 Indicators

Quite standard indicators are used for the comparison of performance of industries:

- welfare indicators: market price, average gross profits and concentration of capital;
- technical efficiency indicators: maximal productivity.

Only the concentration of capital needs some explanation:

$$K = \frac{\sum_j K_j^2}{\left(\sum_j K_j\right)^2}, \tag{11}$$

where K_j is the capital stock of firm j. This indicator gives an equivalent number of firms as if each of them had the same part of capital stock. We have $1 \leq K \leq N$ where N is the number of active firms in the industry. The higher is this indicator, the more evenly balanced is the distribution of capital stock between firms. This is an application of the Herfindall index to the capital stocks and summarizes the inequalities in the distribution of the capital stock.

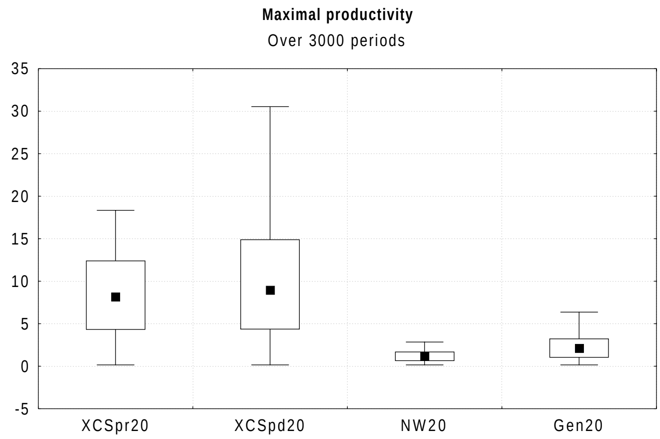

Figure 2: Relative technological efficiency of XCSFirms

3.3 Comparison of industry performance between homogenous cases

Three different points of view can be adopted for the evaluation of the impact of learning on industrial efficiency: technological performance, firms' profit and consumers' welfare. We do not have a direct indicator for consumers' welfare, but the market price is of course inversely related to consumers' surplus. The effect on firms' surplus can be appreciated by comparing the distribution of average cumulative profits in each industry configuration. Technological performance is evaluated through the maximal productivity. The latter shows how far a particular industry can go in the technology space.

The presence of learning firms should normally increase technological efficiency because these firms are able to exploit the increasing relationship that exists between R&D investment and innovation. But there is a specific cost for learning. The overall effect can only assessed through the comparison of different industries.

We compare four homogenous industry configurations with 20 firms:

NW20: with 20 NWFirms;

Gen20: with 20 GenFirms;

XCSpr20: with 20 XCSFirms (objective: profit rate);

XCSpd20: with 20 XCSFirms (objective: productivity).

All configurations have a total population of 20 firms and all GenFirms are the simplest kind, they have $C = 20$ chromosomes of $G = 6$ genes. NWFirms invest 10% in R&D. $rd_{\text{min}} = 3\%$. XCSFirms explore a strategy space composed of 50 actions. Other parameters are common to all simulations and they are given in appendix.

The results on technological efficiency are represented in Figure 2. This graphic shows that the main result of Yildizoglu [2001] is confirmed: industries formed by learning firms exhibit higher technological advancement. In conformity with our preceding results, GenFirms are more efficient than NWFirms but the most significant difference appears when NWFirms are compared to XCSFirms. Moreover this effect is even

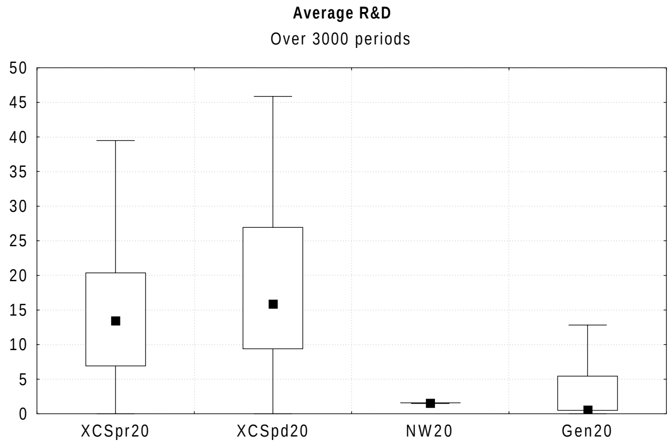

Figure 3: R&D expenses

stronger when the objective of the XCSFirms is their productivity $(XCSpd20)$. This is summarized in the following proposition.

Proposition 1 *Presence of learning firms is a source of technological efficiency at the industry level. The learning process of the XCSFirms is the most efficient one from the point of view of technological advancement.*

The R&D expenses are behind the technological efficiency of learning firms. Figure 3 shows that these firms are able to place themself on a better growth trajectory that allow and imply higher R&D expenses. Yildizoglu [2001] shows that this does not depend on the investment rate of NWFirms. Even with a higher R&D rate, their incapacity to adapt their strategies to the evolution of the industry strongly bounds their growth trajectory and, hence, their R&D expenses. Quite interestingly, this efficiency is even quite costless

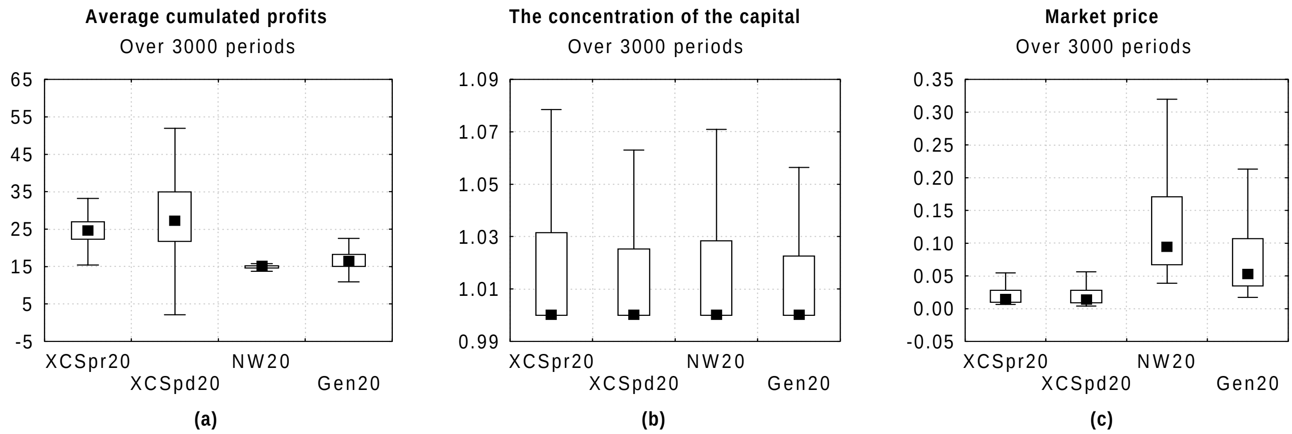

Figure 4: Learning and market structure

for society. Learning of firms does not have a (statistically) significant effect on the concentration of capital (equivalent number of firms in Figure 4–(b)). But it is nevertheless a source of better performance in terms of profit (Figure 4–(a)). Figure 4–(c) clearly shows that these profits do not play against the consumers' welfare: the lower market price implies a higher surplus for the consumers. The global effect of learning on society is clearly positive.

Proposition 2 *Presence of learning firms implies*

1. *higher gross profits;*
2. *lower market price;*
3. *higher social welfare.*

This again confirms our preceding results. These consequences clearly result from the evolution of the arbitrage of learning firms between R&D investment and capital investment. XCSFirms are much more performant in this arbitrage and the use of LCSs clearly implies a better representation of learning. Learning firms deliberately modify both components of their production: cost and capital stock. The presence of learning firms is hence a source of dynamic social efficiency at the industry level and the efficiency at the technological level is the real source of this positive effect on social welfare.

It is interesting to note that XCSpd20 exhibits a higher efficiency than XCSpr20, especially for the profits. This can seem contradictory given that the firms in XCSpr20 are given the profit rate as the objective and one could expect that they get higher profits. This apparent paradox emphasizes the importance of the choice of the reward function in evolutionary learning processes. The productivity of the firms is much better related to their history than their profit rate that is subject to strong fluctuations. This is an important point and we will reconsider it now.

3.4 Learning of firms

In order to better qualify the learning of different firms, we use the data from the last simulation of a batch with longer history (10000 periods) in each configuration.

The Figure 5 shows the R&D rates of firms in three homogenous configurations over the long term. In all configurations, the last part of the history is characterized by the survival of a unique firm. The beginning of the history corresponds to a very high diversity of all R&D rates in $[0,1]$. What is very interesting here, from the point of view of the coherence of the learning process, is the behavior of the surviving firm that obtains the monopoly of the industry.

In each configuration, the monopolist is at the end in quite a stable environment. The only uncertainty in the evolution of its performance is the technological uncertainty (the output of its R&D process): all competitive uncertainty that could results from the decisions of its competitors disappears when it obtains the monopoly of the industry.

This observation explains the behavioral difference we observe in the Figure 5 between *XCSpr* firms (profit–rate maximizers) and *XCSpd* firms (productivity maximizers). As a matter of fact, once the competitive uncertainty is eliminated, the profit–rate maximizers must learn a more complex (indirect and ambiguous) relationship between the profit rate and the R&D rate, while the productivity maximizers learn a direct relationship between the R&D investment and the productivity.

In fact the profit rate of the firm is not directly connected to the R&D rate, but the through the capital stock and the productivity, as it is shown in (12):

$$\begin{aligned}
\pi &= \left(p - \frac{c}{A}\right) \\
&= \left(\frac{\mathbf{D}}{A(rd)K(rd)} - \frac{c}{A(rd)}\right) \\
\frac{\partial\pi}{\partial rd} &= \underbrace{\frac{\partial\pi}{\partial K}}_{\leq 0}\underbrace{\frac{\partial K}{\partial rd}}_{\leq 0} + \underbrace{\frac{\partial\pi}{\partial A}}_{\leq 0}\underbrace{\frac{\partial A}{\partial rd}}_{\geq 0}
\end{aligned} \tag{12}$$

The sign of $\partial K/\partial rd$ is determined by the fact that the capital investment is complementary to the R&D investment (10). The sign of $\partial A/\partial rd$ is defined by the innovation process (5). The global influence is ambiguous and depends on the relative importance of each channel: the locus on the demand curve and the past success of the innovation process. This reward function is consequently very demanding for the learning process. It also strongly depends on the market structure before the establishment of the monopoly.

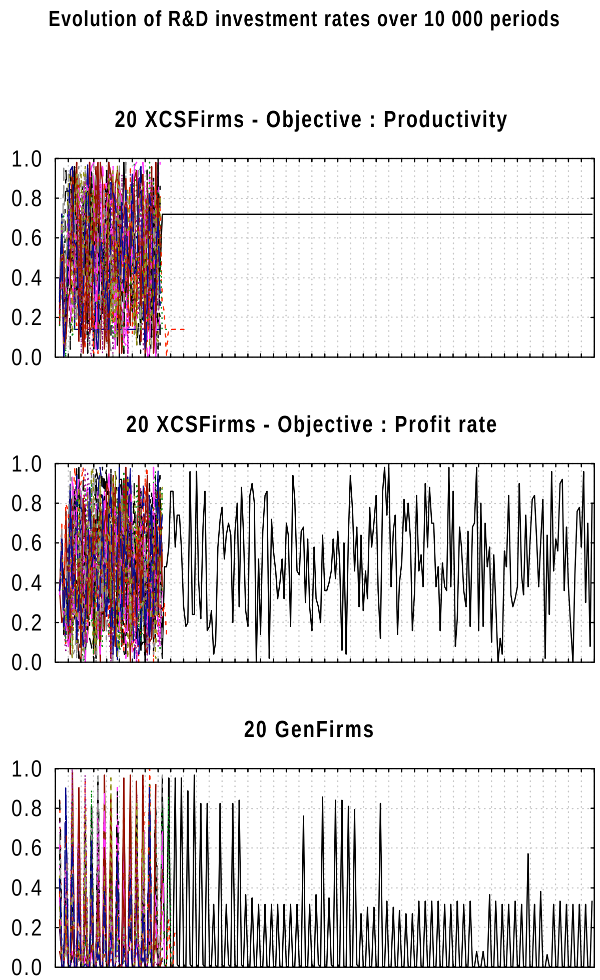

Figure 5: R&D rates from the last simulation

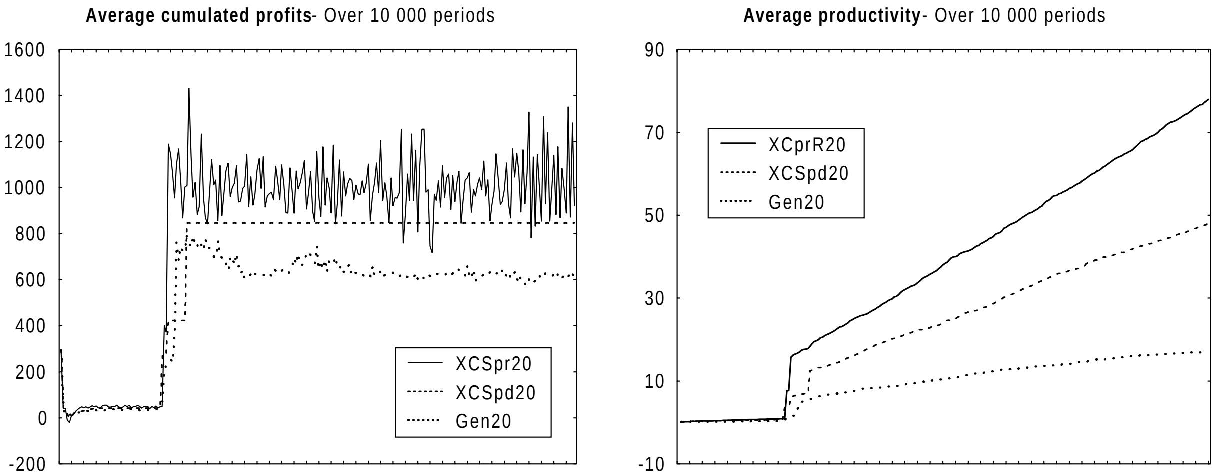

Figure 6: Average profit and productivity in the last simulation

The relation between productivity is much more direct (even if it is random) and unambiguous (see (5)).

It is then clear that the relationship that the productivity maximizers learn remains stable all along of the history, while the relationship between the profit rate and the R&D rate is ambiguous and strongly depends on the market structure. Moreover, given that all market structures but the monopoly are strongly non-Markov, the relative stability of productivity gives a more coherent indication for the learning process.

As a consequence, the learning process of the productivity maximizer monopoly gets stabilized very quickly, while the process of the profit–rate maximizer must take into account the fluctuations of the profit rate.

The Figure 6 clearly exhibits better performance for the profit maximizers. In terms of profit, this better performance establishes the coherence of XCS as a model of learning. The better technological performance of profit maximizers (in comparison with productivity maximizers) must be explained by the fact that the amount of R&D investment depends on the level of the profit as well as on *rd* (see (9).

Both type of firms perform better than the GenFirms. The use of XCS corresponds a more efficient learning model than the simple implementation of GA used for the GenFirms. This strongly advocates for the use of LCSs for modeling learning.

4 Conclusion

This article builds upon the conclusions of Yildizoglu [2001]:

> On a methodological level, one of the shortcomings of Genetic Algorithms in industrial context with endogenous payoff structure is the necessity of effectively using each rule in order to discover its fitness. Learning is consequently slow.

To offset this shortcoming we consider LCSs as a tool for evolutionary modeling of firms' learning. LCSs have the very interesting capacity of generalizing from the past experience of firms. The particular implementation we choose (XCS) has another interesting property: their favour the accuracy of strategies instead of their brute strength. This enhances the expectation dimension of LCSs. This is a very important dimension for modeling learning because many evolutionary models neglect the formation of expectations by the firms and their influence on decisions (see Oltra & Yildizoglu [1999]).

Our results show that the use of XCS pays in terms of the efficiency of the learning process even if they are quite demanding in computational power. Industries formed by XCSFirms are more efficient at the level of technology dynamics, as well as of social welfare. These results clearly advocates for the use of XCS instead of GA for representing learning.

References

Butz, M. V. [2000], XCSJava 1.0: An Implementation of the XCS classifier system in Java , Technical Report 2000027, Illinois Genetic Algorithms Laboratory.

Butz, M. V. & Wilson, S. W. [2000], An Algorithmic Description of XCS, Technical Report 2000017, Illinois Genetic Algorithms Laboratory.

Goldberg, D. E. [1991], *Genetic Algorithms*, Addison-Wesley, Reading: MA.

Jonard, N. & Yildizoglu, M. [1998], 'Technological diversity in an evolutionary industry model with localized learning and network externalities', *Structural Change and Economic Dynamics* **9(1)**, 35–55.

Knight, F. H. [1921], *Risk, Uncertainty and Profits*, number Reprint, Chicago University Press, Chicago.

Kwasnicki, W. & Kwasnicka, H. [1992], 'Market, innovation, competition. an evolutionary model of industrial dynamics', *Journal of Economic Behavior and Organization* **19**, 343–368.

Lanzi, P. L., Stoltzmann, W. & Wilson, S. W. [2000], *Learning Classifier Systems. From Foundations to Applications*, Vol. 1813 of *LNAI*, Springer, Berlin.

Nelson, R. R. & Winter, S. [1982], *An Evolutionary Theory of Economic Change*, The Belknap Press of Harvard University, London.

Oltra, V. & Yildizoglu, M. [1999], 'Expectations and adaptive behaviour: the missing trade-off in models of innovation', *WP BETA No. 9905, Universite Louis Pasteur, Strasbourg* .

Silverberg, G., Dosi, G. & Orsenigo, L. [1988], 'Innovation, diversity and diffusion: a self-organization model', *Economic Journal* **98**, 1032–1054.

Silverberg, J. & Verspagen, B. [1996], From the artificial to the endogenous, *in* M. Helmstadter, Ernst; Perlman, ed., 'Behavioral norms, technological progress, and economic dynamics: Studies in Schumpeterian economics', University of Michigan Press, Ann Arbor.

Simon, H. A. [1958], The role of expectations in adaptive or behavoristic model, *in* M. Bowman, ed., 'Expectations, Uncertainty and Business Behavior', Social Science Council, New York, pp. 49–58.

Simon, H. A. [1976], From substantial to procedural rationality, in Latsis, S. J. (ed), Method and Appraisal in Economics, Cambridge University Press, Cambridge, pp. 129–148.

Watson, C. J., Billingsley, D. J., Croft, D. J. & Huntsberger, D. V. [1993], *Statistics for Management and Economics*, fifth edition, Allyn and Bacon, Boston.

Wilson, S. W. [1995], 'Classifier Fitness Based on Accuracy', *Evolutionary Computation* **3**(2), 149–175. http://prediction-dynamics.com/.

Yildizoglu, M. [2001], 'Competing r&d strategies in an evolutionary industry model', *forthcoming in Computational Economics. Available at http://yildizoglu.montesquieu.u-bordeaux.fr/* .

Appendix: parameter values

d_{in} is fixed in order to have a initial innovation probability of 5%. $d_{im} = d_{in}/4$.

Parameter	Value
Output frequency:	40
Number of simulations:	20
Number of periods: T	2000 or 10000
Using cost of capital: c	0.1
Initial productivity: A_0	0.16
Initial capital: K_0	50
Demand elasticity: η	1
Autonomous demand: **D**	3000
Depreciation rate: δ	1%
Threshold capital: $\underline{K}$	10^{-5}
R&D rate of NWFirms: rd_{NW}	10%
Minimal R&D rate: $rd_{\min}$	3%
Dispersion of Innovations: σ	0.05

Table A1: Common parameters

Parameter	Value
Number of actions:	50
maxPopSize:	800
α	0.1
β	0.3
γ	0.95
Δ	0.1
θ_{GA}	15
θ_{del}	20
θ_{sub}	20
$P[Xover]$	0.8
$P[Mutate]$	0.04
$P[\#]$	0.5
predictionIni	10
fitnessIni	1
Other default values of XCSConstants	

Table A2: Parameters of XCSFirms

Parameter	Value
Number of chromosomes: C	20
Number of genes: G	6
$Xover\ points$	1
Learning rate: n	3
$P[Xover]$	0.7
$P[Mutate]$	0.01

Table A3: Parameters for the GenFirms